arXiv:1402.4844v1 [cs.LG] 19 Feb 2014

The Sample Complexity of Subspace Learning with Partial Information

Alon Gonen* Dan Rosenbaum† Yonina Eldar‡ Shai Shalev-Shwartz§

February 21, 2014

Abstract

The goal of subspace learning is to find a k-dimensional subspace of $\mathbb{R}^d$, such that the expected squared distance between instance vectors and the subspace is as small as possible. In this paper we study the sample complexity of subspace learning in a *partial information* setting, in which the learner can only observe $r \leqslant d$ attributes from each instance vector. We derive upper and lower bounds on the sample complexity in different scenarios. In particular, our upper bounds involve a generalization of vector sampling techniques, which are often used in bandit problems, to matrices.

1 Introduction

Subspace learning is used as a dimensionality reduction technique in a variety of applications such as face recognition, image compression, document and signal classification. In some cases, the access to instances is limited by the nature of the problem. For example, if measurements correspond to the outcome of medical tests, then it is relatively easy to obtain few attributes from each patient but it is not reasonable to obtain all measurements. In this paper we study subspace learning in such partial information setting, and in particular show how to compensate for the lack of information on each individual example by having more examples overall.

The subspace learning problem is formally defined as follow. Let $\mathcal{X}$ be a subset of $\mathbb{R}^d$ and let P be some unknown distribution over $\mathcal{X}$. Our goal is to find a rank-k projection matrix $\Pi \in \mathbb{R}^{d\times d}$, which maps every point $x \in \mathbb{R}^d$ into its closest

*School of Computer Science and Engineering, The Hebrew University, Jerusalem, Israel.
†School of Computer Science and Engineering, The Hebrew University, Jerusalem, Israel
‡Department of Electrical Engineering, Technion, Israel Institute of Technology, Haifa, Israel
§School of Computer Science and Engineering, The Hebrew University, Jerusalem, Israel

point in the k-dimensional subspace corresponding to the range of Π, such that the expected squared distance, $\mathbb{E}_{x\sim P}[\|x - \Pi x\|_2^2]$, is as small as possible.

When $\mathcal{X}$ is a finite set and P is the uniform distribution over $\mathcal{X}$, the optimal solution to the subspace learning problem is the Principal Component Analysis (PCA) solution. Here, we are interested in the more general stochastic optimization setting of subspace learning, in which $\mathcal{X}$ is not restricted to be a finite set and P is an arbitrary distribution over $\mathcal{X}$.

We assume that the learner can obtain i.i.d. sample vectors from $\mathcal{X}$ according to P. In the usual full information setting of subspace learning, the learner has access to all attributes in the sampled vectors. In this paper, we consider subspace learning in the partial information case, in which only r attributes from each vector can be observed. The learner is allowed to choose which of the attributes are revealed. We study the sample complexity of this learning problem, namely, the number of samples that are needed in order to find a projection matrix Π with expected squared distance of at most ϵ more than the optimal expected squared distance. The sample complexity is analyzed in terms of ϵ as well as the dimension d, the dimension of the subspace k, the number of observed attributes r, and properties of the set $\mathcal{X}$. In particular, we focus on sets of the form $\mathcal{X} = \{x \in \mathbb{R}^d : \|x\|_2^2 \leqslant G \wedge \|x\|_\infty \leqslant 1\}$, and study the sample complexity in terms of the scalar G.

We propose three computationally efficient algorithms for subspace learning, named Bandit PCA, Matrix Bandit Gradient Descent (MBGD), and Matrix Bandit Exponentiated Gradient (MBEG). All of these algorithms employ a stochastic strategy for selecting r attributes of each instance. While Bandit PCA and MBGD employ uniform sampling, MBEG uses non-uniform sampling. We derive upper bounds on the sample complexity of these algorithms. The bound for MBGD is $\frac{4d^2G^2k}{r^2\epsilon^2}$. The bound for Bandit PCA is inferior to the bound for MBGD, and we do not know if it is really inferior or that our analysis of Bandit PCA is not tight. For MBEG, we obtain an improved bound for the regime of $r = 2, k = O(1)$, and $\epsilon = O(G^2/d)$. In particular, in this case our bound is $\tilde{O}\left(\frac{dG^2}{\epsilon^2}\right)$, which is smaller than the bound for MBGD by a factor of d.

Next, we derive three lower bounds. The first shows that subspace learning is impossible when only a single attribute of each vector is observed (that is, $r = 1$). The second shows that for the regime $G, k \in O(1)$ the sample complexity is bounded below by $\Omega\left(\frac{d}{r\epsilon}\right)$. We derive a third lower bound, of $\Omega\left(\frac{d^2}{r^2\epsilon}\right)$, which only holds for a family of algorithms, that includes the Bandit PCA, MBGD, and MBEG algorithms mentioned before. This lower bound shows that there is a delicate balance in the dependence of the sample complexity on the different parameters of the problem. In particular, for $r = 2, k = 1$, if ϵ is larger than G^2/d then we need at least d^2/ϵ examples, but if ϵ is smaller than G^2/d then order of

$\frac{dG^2}{\epsilon^2}$ examples suffice. It follows that for small ϵ, the MBEG algorithm is nearly optimal (up to constants and log factors) in its dependence on d and G. Finally, we mention that while the focus of this paper is the stochastic batch learning setting, all of our bounds also hold in the online adversarial learning model.

1.1 Related Work

In the full information case, it has been shown by [16], and [4] that the sample complexity of subspace learning is upper bounded by $O(G^2k/\epsilon^2)$, and this upper bound is achievable by applying PCA on i.i.d. samples according to the distribution P. To the best of our knowledge, no lower bound has been given before, and in Appendix C, we provide a matching lower bound.

Subspace learning in the partial information setting has been studied in [7], where an algorithm named PETRELS is proposed for the problem. However, no formal guarantees are derived for this algorithm.

One possible way to tackle the subspace learning problem with partial information is based on matrix completion techniques. For example, we can think of the partially observed examples as a data matrix with unobserved entries. Then, one could first fill in the missing entries using a matrix completion technique (e.g., the one described in [5]), and then apply PCA on the data matrix. While this approach may work under strong assumptions on the data, our focus in this paper is on methods that come with distribution free guarantees. In fact, our lower bounds provably show that the matrix completion technique would fail in the general case.

Another way to treat the subspace learning problem is by first estimating the data correlation matrix based on the partially observed data and then applying PCA. This is exactly what our Bandit PCA does, and our analysis of this algorithm relies on measure concentration inequalities for matrices, as given in [14].

Our setting resembles the setting of the multi-armed bandit problem ([3]), in which the learner obtains at each time a limited feedback, namely, it receives only the reward of the chosen arm. More concretely, several recent papers studied learning linear predictors with partially observed attributes (e.g. [6]). One can think of our MBGD algorithm as an adaptation of the algorithm given in [6] from vectors to matrices. The main technical difficulty stems from the fact that on the one hand, our examples are d dimensional vectors, and therefore one may think that we need a factor of d/r more examples to compensate for the missing information. On the other hand, the relevant information is in the correlation matrix, whose size is d^2, and therefore maybe we need a factor of $(d/r)^2$ more examples to compensate for the missing information. We show that learning matrices is significantly different by proving that it is impossible to learn with a single observed attribute from each example, because in some cases, a single attribute gives us no information on

the correlation with other attributes, and therefore provides no information on the underlying subspace.

Our MBEG algorithm builds on the matrix Exponentiated Gradient (EG) algorithm of [19]. See also [2, 13]. The main technical novelty here is how to control the exploration-exploitation tradeoff when learning matrices.

2 Main Results

The task of subspace learning in the partial (bandit) information setting is formalized as follows. Let $\mathcal{X} := \{x \in \mathbb{R}^d : \|x\|_2^2 \leqslant G, \|x\|_\infty^2 \leqslant 1\}$. The class of projection matrices from $\mathbb{R}^d$ onto $\mathbb{R}^k$ ($k < d$) is denoted by $\mathcal{P}_k^d$. Given a distribution P over $\mathcal{X}$, the expected loss of a projection matrix Π is defined by $L(\Pi) = \mathbb{E}_{x\sim P}\left[\|x - \Pi x\|^2\right]$.

A subspace learner, $\mathcal{A}$, has access to an oracle that receives r indices, $(i_1, \ldots, i_r) \in [d]^r$, and returns the values of $x_{i_1}, \ldots, x_{i_r}$, for some x which is picked at random according to P. The learner can call the oracle m times and should output a projection matrix, Π.

Definition 1 (Sample complexity) *The* sample complexity *of subspace learning is the function $m : \mathbb{N}^3 \times \mathbb{R}^2 \to \mathbb{N}$, with the property that for every domain $\mathcal{X}$ defined by the parameters d, G, for every subspace dimension k, for every attribute budget r, and for every target accuracy $\epsilon > 0$, $m(d, k, r, G, \epsilon)$ is the minimal integer such that there exists an algorithm $\mathcal{A}$ for which the following holds: For any distribution P over $\mathcal{X}$, if $\mathcal{A}$ calls the oracle $m \geqslant m(d, k, r, G, \epsilon)$ times then it returns a projection matrix Π which satisfies*

$$\mathbb{E}\, L(\Pi) \leqslant \min_{\Pi' \in \mathcal{P}_k^d} L(\Pi') + \epsilon\ ,$$

where the expectation is over the randomness of both the oracle and the algorithm.

The following theorem summarizes our main results. The details and the proofs will follow in the next sections.

Theorem 1 *Assume that $d \geqslant 2$, $k \leqslant \sqrt{d}$, and $r \leqslant d$. The sample complexity of bandit subspace learning is bounded as follows:*

1. *Using Bandit-PCA (Algorithm 1), we obtain*

$$m(d, k, r, G, \epsilon) \leqslant \tilde{O}\left(\frac{d^2 G^3 k^2}{r^2 \epsilon^2}\right)\ .$$

2. *Using Matrix Bandit Gradient Descent (MBGD) (Algorithm 2), we obtain*

$$m(d,k,r,G,\epsilon) \leqslant \frac{4d^2G^2k}{r^2\epsilon^2} \ .$$

3. *Using Matrix Bandit Exponentiated Gradient (MBEG) (Algorithm 3), we obtain*

$$m(d,k,r=2,G,\epsilon) \leqslant \max\left\{\frac{4dG^2k^2\log(d/k)}{\epsilon^2} \ , \ \frac{d^3\log(d/k)}{G^2}\right\} \ .$$

4. *If the learner is limited to (choose and) observe a single attribute of each instance, namely if $r = 1$, then the subspace learning problem is not learnable, that is,*

$$m(d,k,r=1,G,\epsilon) = \infty \ .$$

5. *There exists a family of subspace learners which includes Bandit PCA, MBEG, and MBGD, for which*

$$m(d,k=1,r,G=1,\epsilon) \geqslant \Omega\left(\frac{d^2}{r^2\epsilon}\right) \ .$$

6. *The following lower bound holds for all subspace learners:*

$$m(d,k=1,r,G=1,\epsilon) \geqslant \Omega\left(\frac{d}{r\epsilon}\right) \ .$$

3 Bandit Principal Component Analysis (Bandit PCA)

In this section we describe and analyze the Bandit PCA algorithm, which will prove part 1 of Theorem 1.

The loss of a projection matrix can be rewritten as

$$L(\Pi) = \mathbb{E}\|x - \Pi x\|_2^2 = \mathbb{E}\left[\|x\|_2^2 - 2x^\top\Pi x + x^\top\Pi^\top\Pi x\right] = \mathbb{E}\left[\|x\|_2^2 - x^\top\Pi x\right] \ ,$$

where we used the fact that for every projection matrix $\Pi^\top\Pi = \Pi$. Denoting the inner product of matrices A, B by $\langle A, B\rangle = \mathrm{tr}(AB)$, we can further rewrite the loss as

$$L(\Pi) = \mathbb{E}\left[\|x\|_2^2 - \langle\Pi, xx^\top\rangle\right] = \mathbb{E}\left[\|x\|_2^2\right] - \mathbb{E}\left[\langle\Pi, xx^\top\rangle\right] = \mathbb{E}\left[\|x\|_2^2\right] - \langle\Pi, \mathbb{E}[xx^\top]\rangle \ .$$

Since $\|x\|_2^2$ doesn't depend on Π, the goal of subspace learning is equivalent to finding a projection matrix Π such that

$$\langle \Pi, \mathbb{E}[xx^\top] \rangle \geqslant \max_{\Pi' \in \mathcal{P}_k^d} \langle \Pi', \mathbb{E}[xx^\top] \rangle - \epsilon \,. \tag{1}$$

Denote the correlation matrix by $C = \mathbb{E}[xx^\top]$. It is easy to verify that the matrix which maximizes $\max_{\Pi' \in \mathcal{P}_k^d} \langle \Pi', C \rangle$ is the matrix $\sum_{i=1}^k v_i v_i^\top$, where $v_1, \ldots, v_k$ are the leading eigenvectors of C. The PCA approach for subspace learning in the full information case replaces C with $C(S) = \frac{1}{m} \sum_{i=1}^m x_i x_i^\top$, where $S = (x_1, \ldots, x_m)$ is an i.i.d. sample of examples according to the distribution P. Clearly, $\mathbb{E}[C(S)] = C$. In our case, we will construct an unbiased estimate of C based on partially observed examples, as follows.

Assume that r is even. Pick indices $(i_1, \ldots, i_r)$ uniformly at random. Recall that the learner has access to an oracle which outputs $x_{i_1}, \ldots, x_{i_r}$, for some random $x \sim P$. Denote

$$\hat{x} = \frac{2d}{r} \sum_{s=1}^{r/2} x_{i_s} e_{i_s} \quad \text{and} \quad \hat{y} = \frac{2d}{r} \sum_{t=r/2+1}^{r} x_{i_t} e_{i_t} \,,$$

where e_{i_s} is the all zeros vector except 1 in the i_s coordinate. Then,

$$\mathbb{E}[\hat{x}\hat{y}^\top] = \mathbb{E}\left[\left(\frac{2d}{r} \sum_{s=1}^{r/2} x_{i_s} e_{i_s} \right) \left(\frac{2d}{r} \sum_{q=r/2+1}^{r} x_{i_q} e_{i_q} \right)^\top \right] = \frac{4d^2}{r^2} \sum_{s,q} \mathbb{E}[x_{i_s} x_{i_q} e_{i_s} e_{i_q}^\top]$$

$$= \frac{4d^2}{r^2} \sum_{s,q} \mathbb{E}[C_{i_s,i_q}\, e_{i_s} e_{i_q}^\top] = \frac{4d^2}{r^2} \sum_{s,q} \sum_{i,j=1}^{d} \frac{1}{d^2} C_{i,j}\, e_i e_j^\top = C \,.$$

By averaging $\hat{x}\hat{y}^\top$ over m calls to the oracle, we obtain the Bandit PCA algorithm—see Algorithm 1. Note that we apply the eigendecomposition to a symmetrization of the matrix.

The success of this method depends on how well the average discussed above approximates C. To analyze the approximation, we rely on the following concentration of measure result due to [14]. Below, we use $\|A\|_{\mathrm{sp}}$ to denote the spectral (operator) norm of a matrix A.

Theorem 2 ***(essentially [14][Theorem 3.1])*** *Let C be a matrix such that $\|C\|_{\mathrm{sp}} \leqslant G$. For every $i \in [m]$, let $\hat{C}_i$ be a random matrix of rank 1 which satisfies $\mathbb{E}[\hat{C}_i] = C$, and $\|\hat{C}_i\|_{\mathrm{sp}} \leqslant \gamma$ almost surely. Set $m = \Omega\left(\frac{\gamma G^2 \log(\gamma G^2/\tilde{\epsilon}^2)}{\tilde{\epsilon}^2}\right)$, then for every*

Algorithm 1 Bandit PCA

input: $G, k, d \in \mathbb{N}$, $r \in 2\,\mathbb{N}$
$\hat{C} = 0 \in \mathbb{R}^{d \times d}$
for $i = 1$ **to** m **do**
  pick r indices $i_1, \ldots, i_r \in [d]$, independently and uniformly at random
  let $(x_{i_1}, \ldots, x_{i_r})$ be the output of Oracle$(i_1, \ldots, i_r)$
  $\hat{x} = \frac{2d}{r} \sum_{s=1}^{r/2} x_{i_s} e_{i_s}$
  $\hat{y} = \frac{2d}{r} \sum_{q=r/2+1}^{r} x_{i_q} e_{i_q}$
  $\hat{C}_i = \frac{1}{2} \hat{x} \hat{y}^\top$
  $\hat{C} = \hat{C} + \frac{1}{m} \hat{C}_i$
end for
compute the eigendecomposition $\frac{1}{2}\hat{C} + \frac{1}{2}\hat{C}^\top = \sum_{j=1}^{d} \lambda_j v_j v_j^\top$
Assuming $\lambda_1 \geqslant \ldots \geqslant \lambda_d$, return $\hat{\Pi} = \sum_{j=1}^{k} v_j v_j^\top$

$\tilde{\epsilon} > 0$,

$$P\left(\left\|\frac{1}{m}\sum_{i=1}^{m} \hat{C}_i - C\right\|_{\mathrm{sp}} > \tilde{\epsilon}\right) \leqslant 2\exp\left(-\frac{m\tilde{\epsilon}^2}{\gamma G^2 \log m}\right)$$

Theorem 2 can be translated now into an upper bound on the sample complexity of subspace learning using Bandit PCA (i.e., the bound mentioned in part 1 of Theorem 1). The complete proof is detailed in Appendix A.

4 Matrix Bandit Gradient Descent (MBGD)

We now describe the Matrix Bandit Gradient Descent (MBGD) algorithm. Its analysis proves part 2 of Theorem 1.

Matrix Bandit Gradient Descent is an instance of the Stochastic Gradient Descent (SGD) algorithm. Recall that our problem is equivalent to approximately solving $\min_{\Pi \in \mathcal{P}_k^d} \langle \Pi, -C \rangle$, where $C = \mathbb{E}[xx^\top]$. Our first step is to convexify the optimization problem by replacing the set $\mathcal{P}_k^d$ with its convex hull, $\mathcal{C}_k^d := \mathrm{conv}(\mathcal{P}_k^d)$. Since the objective of the optimization problem is a linear function, for every W in $\mathcal{C}_k^d$, the gradient of the objective is simply $-C$. Therefore, the gradient descent procedure would start with some $W_1 \in \mathcal{C}_k^d$, and at iteration i would update $W_{i+1} = W_i + \eta C$, where $\eta > 0$ is a parameter of the algorithm. Stochastic Gradient Descent allows us to replace C with any unbiased estimate of C. In particular, at iteration i, we construct an unbiased estimator of C by $\hat{C}_i = \frac{1}{2}\hat{x}\hat{y}^\top + \frac{1}{2}\hat{y}\hat{x}^\top$,

where $\hat{x}, \hat{y}$ are as defined in the previous section. We then perform an update of the form $W_{i+1} = W_i + \eta \hat{C}_i$.

The standard implementation of SGD requires us to project W_{i+1} onto the set $\mathcal{C}_k^d$ at the end of each iteration. Since the loss is linear, a lazy implementation of SGD, in which we project only at the last matrix W_{m+1}, yields the same convergence rate as the vanilla SGD (see for example, [15, Corollary 2.17]). The projection step at the end involves the eigendecomposition of W_{m+1}, and the projection of the vector of eigenvalues onto the probabilistic simplex (scaled by k). An efficient implementation of this step is derived, essentially, using [8] and [2]. Denote the projection of W_{m+1} onto $\mathcal{C}_k^d$ by $\hat{W}$. Observe that $\hat{W}$ is not necessarily a member of $\mathcal{P}_k^d$. Therefore, the final step of the algorithm is to construct an element from $\mathcal{P}_k^d$ that performs similarly to $\hat{W}$. To do so, we rely on the following lemma, due to [19]:

Lemma 1 ([19]) *Every matrix $W \in \mathcal{C}_k^d$ can be decomposed in time $O(d^3)$ into a convex combination of at most d elements from $\mathcal{P}_k^d$.*

For completeness, we recall the decomposition procedure of [19] in Appendix D. Getting back to our algorithm, let $\hat{W} = \sum_{j=1}^{d} \alpha_j \Pi_j$ be the decomposition of $\hat{W}$ into a convex combination of elements from $\mathcal{P}_k^d$ obtained from the decomposition procedure. The final step of MBGD sets the output matrix, $\hat{\Pi}$, to be Π_j with probability α_j. This guarantees that the expected performance of $\hat{\Pi}$ is the same as the performance of $\hat{W}$. The resulting algorithm is detailed in Algorithm 2.

Remark 1 *Observe that the total runtime of MBGD to find an ϵ-accurate solution is $O\left(\frac{d^2 G^2 k}{\epsilon^2} + d^3\right)$. As a comparison, the runtime of the vanilla PCA for the full information case is the same (assuming that PCA is implemented using a complete eigendecomposition of the empirical correlation matrix). That is, at least in some regimes, we loose nothing in runtime by constraining the algorithm to observe only r attributes from each instance.*

Proof 1 ***(of part 2 of Theorem 1)*** *Denote by $\hat{\Pi}$ the output of MBGD, and denote by $\Pi^\star$ the optimal projection. Recall that $C = \mathbb{E}[xx^\top]$. The standard analysis of SGD ([12]) implies that for every $\Pi^\star \in \mathcal{P}_k^d$,*

$$\mathbb{E}\left[\langle \tilde{W}, -C\rangle\right] - \langle \Pi^\star, -C\rangle \leqslant \frac{\eta}{2} \max_{i\in[m]} \mathbb{E}\left[\|\hat{C}_i\|_F^2\right] + \frac{\|\Pi^\star - W_1\|_F^2}{2\eta m} \,.$$

The term $\|\Pi^\star - W_1\|_F^2$ is at most $2k$. Recall that $\hat{C}_i = \frac{1}{2}\hat{x}\hat{y}^\top + \frac{1}{2}\hat{y}\hat{x}$ for some random x. Since given x, $\hat{x}$ and $\hat{y}$ are independent, we obtain that $\mathbb{E}[\|\hat{x}\hat{y}^\top\|_F^2] =$

Algorithm 2 Matrix Bandit Gradient Descent

input: $G, k, d \in \mathbb{N}$, $r \in 2\,\mathbb{N}$
$\eta = \sqrt{\frac{k}{d^2 G^2 m}}$
$W_1 = \frac{k}{d} I$
for $i = 1$ **to** m **do**
  pick r indices $i_1, \ldots, i_r \in [d]$, independently and uniformly at random
  let $(x_{i_1}, \ldots, x_{i_r})$ be the output of Oracle$(i_1, \ldots, i_r)$
  $\hat{x} = \frac{2d}{r} \sum_{s=1}^{r/2} x_{i_s} e_{i_s} s$
  $\hat{y} = \frac{2d}{r} \sum_{q=r/2+1}^{r} x_{i_q} e_{i_q}$
  $\hat{C}_i = \frac{1}{2} \hat{x} \hat{y}^\top + \frac{1}{2} \hat{y} \hat{x}^\top$
  $W_{i+1} = W_i + \eta \hat{C}_i$
end for
compute eigendecomposition $W_{m+1} = \sum_{j=1}^{d} \lambda_j v_j v_j^\top$ s.t. $\lambda_1 \geqslant \ldots \geqslant \lambda_d$
$\rho = \max \left\{ j \in [d] : \lambda_j - \frac{1}{j} \left(\sum_{s=1}^{j} \lambda_j - k \right) > 0 \right\}$
$\theta = \frac{1}{\rho} \left(\sum_{j=1}^{\rho} \lambda_j - k \right)$
for each $j \in [d]$, set $\tilde{\lambda}_j = \max\{\lambda_j - \theta, 0\}$
$\tilde{W} = \sum_{j=1}^{d} \tilde{\lambda}_j v_j v_j^\top$
decompose $\tilde{W}$ into $\tilde{W} = \sum_{j=1}^{d} \alpha_j \Pi_j$ using Algorithm 4
return $\hat{\Pi} = \Pi_j$ with probability α_j

$\mathbb{E}[\|\hat{x}\|_2^2] \cdot \mathbb{E}[\|\hat{y}\|_2^2]$. *Next, observe that*

$$\mathbb{E}[\|\hat{y}\|_2^2] = \mathbb{E}[\|\hat{x}\|_2^2] = \mathbb{E}\sum_{s=1}^{r/2}\left(\frac{2d}{r}x_{i_s}\right)^2 = \left(\frac{2d}{r}\right)^2\sum_{s=1}^{r/2}\mathbb{E}\,x_{i_s}^2$$
$$= \left(\frac{2d}{r}\right)^2 (r/2)\sum_{j=1}^{d}\frac{1}{d}x_j^2 = \frac{2d}{r}G\ .$$

Hence, $\mathbb{E}[\|\hat{C}_i\|_F^2] \leqslant \frac{4d^2G^2}{r^2}$. *Finally, substitute* $\eta = \sqrt{\frac{k}{d^2G^2m}}$ *and use the fact that given* $\tilde{W}$, $\mathbb{E}\hat{\Pi} = \tilde{W}$, *to conclude the proof.*

5 Matrix Bandit Exponentiated Gradient (MBEG)

We next describe the Matrix Bandit Exponentiated Gradient (MBEG) algorithm. Its analysis concludes the proof of part 3 of Theorem 1. In short, the main differences between MBEG and MBGD are as follows:

1. MBEG uses a more sophisticated sampling method which gives higher priority to "stronger" directions (according to the current matrix W_t).
2. Update rule: while MBGD is derived using the (stochastic) Gradient Descent algorithm, MBEG is derived using the Exponentiated Gradient algorithm for matrices given in [17].

Throughout this section we focus on the case $r = 2$.

5.1 Prioritizing "stronger" directions

The sampling mechanism applied to estimate the gradient in MBGD seems very naive as it does not take into account any information gathered through the learning process. In contrast, the MBEG algorithm samples the oracle with attributes that depend on its current matrix $W \in \mathcal{C}_k^d$. Recall that for every $W \in \mathcal{C}_k^d$ we have $\sum_i W_{i,i} = k$. Therefore, one natural way to pick coordinates is according to the probability $p_{s,q} = \frac{W_{s,s}}{k} \cdot \frac{W_{q,q}}{k}$. Unfortunately, we could not obtain a good bound using this sampling technique, as it yielded a high variance of the estimate. Instead, we picked attributes with probability $p_{s,q}$ defined as:

$$p_{s,q} = (1-\alpha)\frac{W_{s,s} + W_{q,q}}{2dk} + \frac{\alpha}{d^2}\ , \tag{2}$$

for some parameter $\alpha \in (0,1)$, which is tuned below. That is, we mix a uniform distribution over the $[d]^2$ elements, with a distribution which gives higher sampling probability to pairs for which $W_{s,s}, W_{q,q}$ are high, reflecting a bias toward sampling from "stronger" directions. The mixing with the uniform distribution guarantees that every pair has large enough probability to be sampled, which will later help us ensure that we perform enough "exploration".

Based on this probability distribution over pairs, we define an unbiased estimate of the matrix C by

$$\hat{C} = \frac{1}{2p_{s,q}} x_q x_q (E_{s,q} + E_{q,s}) \,, \tag{3}$$

where $E_{s,q}$ is the all zeros matrix except 1 in the i,j coordinate.

5.2 The Exponentiated Gradient update rule

MBEG is derived from the extension of the Exponentiated Gradient (EG) ([10]) algorithm to matrices, which is presented in [17]. The basic idea can be described as follows. Let $R(W) = \mathrm{tr}(W \log W - W)$ be the *quantum entropy* regularizer. This function yields the following Bregman divergence:

$$D_R(W,U) = \mathrm{tr}(W \log W - W \log U - W + U) \,.$$

Given the current matrix W_t, let $\hat{C}_t$ be an unbiased estimate of $C = \mathbb{E}[xx^\top]$ according to Equation (3) and Equation (2). The update rule of MBEG is composed of the following two steps:

1. $U_{t+1} = \exp(\log(W_t) + \eta \hat{C}_t)$
2. $W_{t+1} = \operatorname*{argmin}_{W \in \mathcal{C}_k^d} D_R(W, U_{t+1})$

The resulting algorithm is spelled out in Algorithm 3.

5.3 Analysis of MBEG

We turn to analyze the sample complexity of MBEG. We begin with the next bound, due to [9], which essentially follows from [17].

Theorem 3 ***([9][Theorem 13])***
Assume that the sequence $\hat{C}_1, \ldots, \hat{C}_m$ obtained during the run of MBEG satisfies $\forall i,\ \|\hat{C}_i\|_{\mathrm{sp}} \leqslant 1/\eta$. Then, for every $\Pi^\star \in \mathcal{P}_k^d$,

$$\sum_{i=1}^{m} \langle W_i - \Pi^\star, -\hat{C}_i \rangle \leqslant \frac{k \log d}{\eta} + \eta \sum_{i=1}^{m} \langle W_i, \hat{C}_i^2 \rangle \,. \tag{4}$$

Algorithm 3 Matrix Bandit Exponentiated Gradient

Parameters:
m: number of iterations
$r = 2$ (attribute budget)
$\eta = \sqrt{\frac{\log(d/k)}{dmG^2}}$
$\alpha = \frac{1}{2}\eta d^2$ (we assume m is large enough so that $\alpha \leqslant 1/2$)
Initialize: $W_1 = \frac{k}{d} I$
for $i = 1$ **to** m **do**
  denote $W = W_i$
  pick (s, q) with probability $p_{s,q} = (1-\alpha)\frac{W_{s,s}+W_{q,q}}{2dk} + \alpha\frac{1}{d^2}$
  let (x_s, x_q) be the output of Oracle(s, q)
  $\hat{C}_i = \frac{1}{2p_{s,q}}\, x_s x_q (E_{s,q} + E_{q,s})$
  $U_{i+1} = \exp(\log(W_i) + \eta \hat{C}_i)$
  $W_{i+1} = \operatorname{argmin}_{W \in \mathcal{C}_k^d} D_R(W, U_{i+1})$
end for
$\overline{W} = \frac{1}{m}\sum_{i=1}^m W_i$
decompose $\overline{W}$ into $\overline{W} = \sum_{j=1}^d \alpha_j \Pi_j$ using Algorithm 4
return $\hat{\Pi} = \Pi_j$ with probability α_j

Using the definition of $\hat{C}_i$ we have that $\hat{C}_i^2 = \hat{C}_i \hat{C}_i^\top$ is a diagonal matrix with all elements on the diagonal equals to zero except the (s, s) and (q, q) elements, which both equal to $\frac{x_s^2 x_q^2}{4p_{s,q}^2}$. It follows that

$$\|\hat{C}_i\|_{\mathrm{sp}} = \frac{|x_s x_q|}{2p_{s,q}} \leqslant \frac{1}{2p_{s,q}} \leqslant \frac{d^2}{2\alpha} = \frac{1}{\eta}\,.$$

Therefore the conditions of Theorem 3 hold. Taking expectation over Equation (4), we obtain

$$\sum_{i=1}^m \mathbb{E}\langle W_i - \Pi^\star, -\hat{C}_i\rangle \leqslant \frac{k\log(d/k)}{\eta} + \eta \sum_{i=1}^m \mathbb{E}\langle W_i, \hat{C}_i^2\rangle\,. \tag{5}$$

Let $\mathbb{E}_i$ denote the conditional expectation given W_i. Then, $\mathbb{E}_i \hat{C}_i = C$ and therefore, by the law of total expectation,

$$\mathbb{E}\langle W_i - \Pi^\star, -\hat{C}_i\rangle = \mathbb{E}\langle W_i - \Pi^\star, -C\rangle\,. \tag{6}$$

Next, fix some $i \in [m]$, and let $W = W_i, \hat{C} = \hat{C}_i$. Then,

$$\begin{aligned}\mathbb{E}_i \langle W, \hat{C}^2 \rangle &= \sum_{(s,q)\in[d]^2} p_{s,q}(W_{s,s} + W_{q,q}) \frac{1}{4p_{s,q}^2} x_s^2 x_q^2 \\ &= \sum_{(s,q)\in[d]^2} \frac{(W_{s,s} + W_{q,q})}{4p_{s,q}} x_s^2 x_q^2\end{aligned}$$

By our assumption that $\alpha \in (0, 1/2)$, we obtain that

$$\frac{W_{s,s} + W_{q,q}}{p_{s,q}} = \frac{W_{s,s} + W_{q,q}}{(1-\alpha)\frac{W_{s,s}+W_{q,q}}{2dk} + \alpha/d^2} \leqslant \frac{W_{s,s} + W_{q,q}}{(1-\alpha)\frac{W_{s,s}+W_{q,q}}{2dk}} \leqslant 4dk \ .$$

Therefore

$$\mathbb{E}_i \langle W, \hat{C}^2 \rangle \leqslant dk \sum_{(s,q)\in[d]^2} x_s^2 x_q^2 \leqslant dkG^2 \ . \tag{7}$$

Plugging Equation (6) and Equation (7) into Equation (5) we obtain that,

$$\mathbb{E} \sum_{i=1}^{m} \langle W_i - \Pi^\star, -C \rangle \leqslant \frac{k \log(d/k)}{\eta} + \eta m d G^2 k$$

Dividing by m, denoting $\overline{W} = \frac{1}{m}\sum_{i=1}^m W_i$, substituting $\eta = \sqrt{\frac{\log(d/k)}{dmG^2}}$, rearranging terms, and observing that $\mathbb{E}[\hat{P}|\overline{W}] = \overline{W}$, we have that

$$\langle \mathbb{E}\,\hat{\Pi} - \Pi^\star, -C \rangle = \langle \mathbb{E}\,\overline{W} - \Pi^\star, -C \rangle \leqslant 2\sqrt{\frac{dG^2k^2\log(d/k)}{m}} \ .$$

The right-hand side of the above would be smaller than ϵ if $m \geqslant \frac{4dG^2k^2\log(d/k)}{\epsilon^2}$. Note, however, that we also require that m is large enough so that $\alpha \leqslant 1/2$. Since $\alpha = \frac{\eta d^2}{2}$, it follows that m should also satisfy $m \geqslant \frac{d^3\log(d/k)}{G^2}$. This concludes the proof of part 3 of Theorem 1.

6 Lower Bounds

We now derive three lower bounds on the sample complexity of bandit subspace learning. We are mainly interested in the dependence on the parameter d. Since we are dealing with lower bounds, and are interested in the dependence on d, r, G, ϵ, we will focus on the case $k = 1$. We emphasize the fact that all the lower bounds are proved in the realizable setting, i.e., the expected loss of the optimal projection matrix is zero. The complete proof of the lower bounds are detailed in Appendix B.

6.1 Impossibility of Subpace Learning with a Single Attribute

Part 4 of Theorem 1 implies that subspace learning with a single attribute ($r = 1$) is impossible. The idea of the proof is as follows. For each $s \in [d]$, define $u_s = \sqrt{\frac{G}{d}}(\sum_{j \neq s} e_j - e_s)$, and let P_s be the uniform distribution over $\{u_s, -u_s\}$. Note that for each P_s, there exists a projection matrix (namely, $(u_s/\|u_s\|)(u_s/\|u_s\|)^\top$) whose expected loss is zero. However, for every i, the value of Oracle(i) is distributed uniformly over $\{-1, 1\}$, and therefore gives no information at all to the learner.

6.2 The sample complexity of *dyadic* subspace learners is (at least) quadratic in the dimension

We introduce a family of subspace learners, which we call *dyadic* subspace learners. This family includes the three subspace learners presented in this paper. We will show that the sample complexity of any algorithm in this family is at least $\Omega\left(\frac{d^2}{r^2\epsilon}\right)$, and hence deduce part 5 of Theorem 1.

Before we formally define the family of dyadic algorithms, it would be more instructive to describe an adversarial distribution which will pose a challenge to all algorithms in this family. Let $G = 1$, and let $\epsilon \in (0, 1/4)$. Draw at random an index $s \in [d]$. A distribution P_s over $\mathcal{X}$ is defined such that the probability of the all zeros vector is $1 - \epsilon$, and the probability of the vector e_s is ϵ. As we show in Appendix B, a subspace learner must identify the index s in order to return an ϵ-optimal projection matrix with respect to P_s.

Definition 2 *A subspace learner is called* dyadic *if it satisfies the following conditions:*

- *It maintains a distribution for attribute selection, whose initial value is the uniform distribution over* $[d]^2$.
- *Let* $(x_1, \ldots, x_r) =$ *Oracle*$(i_1, \ldots, i_r)$. *The update of the attribute sampling distribution of the algorithm only depends on* $\{x_i x_j : x_i x_j \neq 0, i \in [r/2], j \in [r] \backslash [r/2]\}$.

Note that the algorithms Bandit PCA, MBGD, and MBEG are all dyadic. Let $\mathcal{A}$ be a dyadic learner. Assuming the distribution P_s defined above, it is not hard to see that whenever the size of the sample is at most $o\left(\frac{d^2}{r^2\epsilon}\right)$, then with non-neglectable probability, the sampling distribution of the learner will remain uniform, hence it will fail to find an ϵ-accurate solution.

6.3 The price of bandit information is at least linear in the dimension

We now prove part 6 of Theorem 1, i.e. that the price of bandit information is at least d. Let $\epsilon \in (0, 1/4)$, and let $\mathcal{A}$ be a bandit subspace learner. For each $s \in [d]$, the distribution P_s is defined as follows. The zero vector is drawn with probability $1 - \epsilon$, and the vector e_s is drawn with probability ϵ. Note that $\mathbb{E}[xx^\top] = \epsilon\, e_s e_s^\top$, and thus the optimal projection is simply $e_s e_s^\top$. If P_s is chosen in advance at random, and the size of the sample is $o\left(\frac{d}{r\epsilon}\right)$, then, with non-neglectable probability, the learner $\mathcal{A}$ observes only zeros, and hence $\mathcal{A}$ cannot distinguish between the different P_s.

7 Discussion and Open Problems

We initiated a study of the sample complexity of subspace learning in a partial information setting. We derived several upper bounds by employing computationally efficient algorithms, and derived several lower bounds. While the sample complexity for the full information case is well understood, the sample complexity in the partial information case is still not fully characterized. Improving the upper and/or the lower bounds is left to future work.

References

[1] Martin Anthony and Peter L Bartlett. *Neural network learning: Theoretical foundations*. cambridge university press, 2009.

[2] Raman Arora, Andrew Cotter, and Nathan Srebro. Stochastic optimization of pca with capped msg. *arXiv preprint arXiv:1307.1674*, 2013.

[3] Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert E Schapire. The nonstochastic multiarmed bandit problem. *SIAM Journal on Computing*, 32(1):48–77, 2002.

[4] Gilles Blanchard, Olivier Bousquet, and Laurent Zwald. Statistical properties of kernel principal component analysis. *Machine Learning*, 66(2-3):259–294, 2007.

[5] Emmanuel J Candès and Benjamin Recht. Exact matrix completion via convex optimization. *Foundations of Computational mathematics*, 9(6):717–772, 2009.

[6] Nicolo Cesa-Bianchi, Shai Shalev-Shwartz, and Ohad Shamir. Efficient learning with partially observed attributes. *The Journal of Machine Learning Research*, 12:2857–2878, 2011.

[7] Yuejie Chi, Yonina C Eldar, and Robert Calderbank. Petrels: Subspace estimation and tracking from partial observations. *IEEE Transactions Signal Processing*, 61:5947–5959, 2013.

[8] John Duchi, Shai Shalev-Shwartz, Yoram Singer, and Tushar Chandra. Efficient projections onto the l 1-ball for learning in high dimensions. In *Proceedings of the 25th international conference on Machine learning*, pages 272–279. ACM, 2008.

[9] Elad Hazan, Satyen Kale, and Shai Shalev-Shwartz. Near-optimal algorithms for online matrix prediction. *arXiv preprint arXiv:1204.0136*, 2012.

[10] Jyrki Kivinen and Manfred K Warmuth. Exponentiated gradient versus gradient descent for linear predictors. *Information and Computation*, 132(1):1–63, 1997.

[11] Wouter M Koolen, Wojciech Kotlowski, and Manfred K Warmuth. Learning eigenvectors for free. In *Advances in Neural Information Processing Systems*, pages 945–953, 2011.

[12] Arkadi Semenovich Nemirovski and David Yudin. *Problem complexity and method efficiency in optimization*. Wiley (Chichester and New York), 1983.

[13] Jiazhong Nie, Wojciech Kotłowski, and Manfred K Warmuth. Online pca with optimal regrets. In *Algorithmic Learning Theory*, pages 98–112. Springer, 2013.

[14] Mark Rudelson and Roman Vershynin. Sampling from large matrices: An approach through geometric functional analysis. *Journal of the ACM (JACM)*, 54(4):21, 2007.

[15] S. Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends® in Machine Learning*, 4(2):107–194, 2011.

[16] John Shawe-Taylor, Christopher KI Williams, Nello Cristianini, and Jaz Kandola. On the eigenspectrum of the gram matrix and the generalization error of kernel-pca. *Information Theory, IEEE Transactions on*, 51(7):2510–2522, 2005.

[17] Koji Tsuda, Gunnar Rätsch, and Manfred K Warmuth. Matrix exponentiated gradient updates for on-line learning and bregman projection. In *Journal of Machine Learning Research*, pages 995–1018, 2005.

[18] Manfred K Warmuth. When is there a free matrix lunch? In *Learning Theory*, pages 630–632. Springer, 2007.

[19] M.K. Warmuth and D. Kuzmin. Randomized online pca algorithms with regret bounds that are logarithmic in the dimension. *Journal of Machine Learning Research*, 9:2287–2320, 2008.

A Proof of the Upper Bound on The Sample Complexity of Bandit PCA

In this section we provide a detailed proof of part 1 of Theorem 1. We use the following lemma.

Lemma 2 *Fix any estimate $\hat{x}\hat{y}^\top$ obtained during the run of Bandit PCA. Then, with probability* 1*,*

$$\|\hat{x}\hat{y}^\top\|_{\mathrm{sp}} \leqslant \frac{4d^2G}{r^2} \,.$$

Proof 2

$$\begin{aligned}
\|\hat{x}\hat{y}^\top\|_{\mathrm{sp}} &= \frac{4d^2}{r^2}\left\|\sum_{s=1}^{r/2}\hat{x}_{i_s}\sum_{q=r/2+1}^{r}\hat{y}_{i_q}^\top\right\|_{\mathrm{sp}} \\
&\leqslant \frac{4d^2}{r^2}\left\|\sum_{s=1}^{r/2}\hat{x}_{i_s}\right\|_2\left\|\sum_{q=r/2+1}^{r}\hat{y}_{i_q}\right\|_2 \\
&\leqslant \frac{4d^2}{r^2}G \,.
\end{aligned}$$

Proof 3 ***(of part 1 of Theorem 1)*** *The proof is divided into two parts. In the first part, relying on Theorem 2, we prove that $\tilde{O}\left(\frac{d^2G^3k^2}{r^2\epsilon^2}\right)$ samples suffice to guarantee that with probability at least $1-\delta$ the averaged correlation matrix obtained during the run of Bandit PCA, denoted $\hat{C}(S) = \frac{1}{m}\sum_{i=1}^{m}\hat{C}_i$, satisfies*

$$\|\hat{C}(S) - C\|_{\mathrm{sp}} \leqslant \epsilon/(2k) \,.$$

In the second part, we prove that when $\hat{C}(S)$ is $\epsilon/(2k)$-close to C, the output of Bandit PCA is ϵ-optimal.

Let P be an unknown distribution over $\mathcal{X}$, and let $m = \tilde{\Omega}\left(\frac{d^2G^3k^2}{r^2\epsilon^2}\right)$. Recall that $C = \mathbb{E}[xx^\top]$. As we mentioned abobe, the condition $\|C\|_{\mathrm{sp}} \leqslant G$ is satisfied by the assumption that $\mathcal{X} \subseteq \{x : \|x\|_2^2 \leqslant G\}$. Also, we mentioned that each of the matrices $\hat{C}_i$ produced during the run of Bandit PCA forms an unbiased estimate of C. Combining Lemma 2 with the triangle inequality implies that for each $i \in [m]$,

$$\|\hat{C}_i\|_{\mathrm{sp}} \leqslant \frac{1}{2}\left(\|\hat{x}_i\hat{y}_i^\top\|_{\mathrm{sp}} + \|\hat{y}_i\hat{x}_i^\top\|_{\mathrm{sp}}\right) \leqslant \frac{4d^2G}{r^2}\ ,$$

almost surely. We now apply Theorem 2, along with the inequality $\ln m \leqslant \alpha m - \ln\alpha - 1$, which holds for every $\alpha, m > 0$ ([1][Inequality 1.2, Appendix 1]). It follows that $\hat{C}(S)$ is $\epsilon/(2k)$-close to C with probability at least $1-\delta$.

Assume now that $\|\hat{C}(S) - C\|_{\mathrm{sp}} \leqslant \epsilon/(2k)$. Recall that the trace norm of a matrix A in $\mathbb{R}^{d\times d}$ is defined by $\|A\|_{\mathrm{tr}} = \sum_{j=1}^d \sigma(A)_j$, where $(\sigma(A)_j)_{j=1}^d$ is the sequence of the singular values of A. Note that the trace norm of each projection matrix $\Pi \in \mathcal{P}_k^d$ is exactly k. It follows by Hölder's inequality that for every $\Pi \in \mathcal{P}_k^d$,

$$\begin{aligned}|\langle \Pi, \hat{C}(S)\rangle - \langle \Pi, C\rangle| &\leqslant \|\Pi\|_{\mathrm{tr}}\,\|\hat{C}(S) - C\|_{\mathrm{sp}} \\ &\leqslant k\cdot\epsilon/(2k) \\ &= \epsilon/2\ .\end{aligned}$$

Combining the last inequality with the fact that $\hat{\Pi}$ maximizes $\langle \Pi, \hat{C}(S)\rangle$, we obtain that for every $\Pi \in \mathcal{P}_k^d$,

$$\langle \Pi, C\rangle - \langle \hat{\Pi}, C\rangle \leqslant \langle \Pi, \hat{C}(S)\rangle + \epsilon/2 - (\langle \hat{\Pi}, \hat{C}(S)\rangle - \epsilon/2) \leqslant \epsilon\ .$$

In view of Equation (1), we conclude that with probability $1-\delta$, Bandit PCA returns an ϵ-optimal projection matrix.

B Proofs of the Lower Bounds on Subspace Learning with Partial Information

In this section we provide detailed proofs of our lower bounds on subspace learning with partial information. The sketches of these proofs were presented in Section 6.

B.1 Proof of part 4 of Theorem 1

Proof 4 *Recall that $k = 1$. We will assume w.l.o.g. that $G = 1$ (the proof for other values of G is obtained from this proof by simple scaling arguments). Recall the definition of the distributions $\{P_s : s \in [d]\}$. For each $s \in [d]$, define*

$u_s = \frac{1}{\sqrt{d}}(\sum_{j \neq s} e_j - e_s)$, *and let* P_s *be the uniform distribution over* $\{u_s, -u_s\}$. *A concrete distribution is set by choosing* $s \in [d]$ *uniformly at random.*

B.1.1 A successful subspace learner must identify the index s

Let $\mathcal{A}$ *be a subspace learner which is allowed to choose and observe a single attribute of each instance. We first prove that* $\mathcal{A}$ *must identify the chosen index* s *in order to return an* ϵ*-optimal projection matrix.*

We fix some $s \in [d]$*, and analyze the subspace learning task accordingly. Note that*

$$\mathbb{E}[xx^\top] = u_s u_s^\top \ .$$

In view of Equation (1), we conclude that the optimal projection matrix is $\Pi^\star = u_s u_s^\top$. *Furthermore, denoting the output of* $\mathcal{A}$ *by* $\hat{\Pi} = \hat{u}\hat{u}^\top$, *where* $\hat{u}$ *is a unit vector, we obtain that*

$$\begin{aligned} L(\hat{\Pi}) - L(\Pi^\star) &= \langle u_s u_s^\top, u_s u_s^\top \rangle - \langle \hat{u}\hat{u}^\top, u_s u_s^\top \rangle \\ &= 1 - \langle \hat{u}\hat{u}^\top, u_s u_s^\top \rangle \ . \end{aligned}$$

Denote $\langle \hat{u}\hat{u}^\top, u_s u_s^\top \rangle$ *by* θ_s*, (in words,* θ_s *"measures" the correlation between* $\hat{u}$ *and* u_s*). Note that if* $L(\hat{\Pi}) - L(\Pi^\star) < 2\epsilon$*, then* $\theta_s > 1 - 2\epsilon$. *The following computation shows that if* $\theta_s > 1 - 1/d^2$*, then for any* $q \in [d] \backslash \{s\}$, $\theta_q \leqslant 1 - 1/d^2$ *(in words, if* $\hat{u}$ *is highly correlative with* u_s*, then* $\hat{u}$ *is not highly correlative with* u_q*).*

$$\begin{aligned} \theta_q &= \theta_q + \theta_s - \theta_s \\ &= \langle \hat{u}\hat{u}^\top, u_q u_q^\top + u_s u_s^\top \rangle - \theta_s \\ &\leqslant \langle \hat{u}\hat{u}^\top, u_q u_q^\top + u_s u_s^\top \rangle - 1 + 2\epsilon \\ &= \|\hat{u}\hat{u}^\top\|_F \|u_q u_q^\top + u_s u_s^\top\|_F - 1 + 2\epsilon \\ &\leqslant 1 \cdot 2\sqrt{\frac{d^2 - 2d + 4}{d^2}} - 1 + 2\epsilon \end{aligned}$$

Letting $\epsilon \in (0, 1/(2d) - 1/d^2)$*, we obtain that* $\theta_q < 1 - 2\epsilon$. *Consider the task of recognizing the distribution* P_s. *With the same assumptions on the environment (and the interaction between the environment and the learner), the role of the learner now is to recognize* P_s*, or equivalently to recognize* s*.From above, we obtain a simple reduction from this task to the task of subspace learning. Namely,*

given the output $\hat{\Pi} = \hat{u}\hat{u}^\top$ of a subspace learner $\mathcal{A}$, return $\hat{s} = \arg\max_{s' \in [d]} \theta_s$. We conclude from above that a successful subspace learner can be used to identify the distribution, i.e., to identify s.

B.1.2 The impossibility of identifying s

As we mentioned in the sketch of this proof, assuming that the learner has access to a single attribute, the learner can not distinguish between the distributions, since any single attribute is distributed uniformly over $\{-1, 1\}$. Hence, for any sample size, and any strategy of the learner, the expected loss of the learner's output is at least $\frac{d-1}{d} 2\epsilon > \epsilon$. The probabilistic method can be applied now to deduce that for any learning algorithm, and any sample size, there exist $\epsilon > 0$, and a distribution P, for which $\mathbb{E}_{S \sim P^m} L(\hat{\Pi}) - L(\Pi^\star) > \epsilon$.

B.2 Proof of part 5 of Theorem 1

Proof 5 *Recall that $k = 1$, and $G = 1$. Let $\epsilon \in (0, 1/4)$. Recall the definition of the distributions $\{P_s : s \in [d]\}$. The probability of the zero vector is $1 - c\epsilon$ (for $c > 2$), and the probability of the vector e_s is $c\epsilon$. A concrete distribution is set by choosing $s \in [d]$ uniformly at random.*

B.2.1 A successful subspace learner must identify the drawn index s

Let $\mathcal{A}$ be a dyadic subspace learner, and let P_s be a concrete distribution. Note that

$$\mathbb{E}[xx^\top] = ce_s e_s^\top = c\epsilon E_{s,s} \ .$$

In view of Equation (1), the optimal projection matrix is $\Pi^\star := E_{s,s}$. Furthermore, denoting the output of $\mathcal{A}$ by $\hat{\Pi} = \hat{u}\hat{u}^\top$, we obtain that if $L(\hat{\Pi}) - L(\Pi^\star) < \epsilon < c\epsilon/2$, then $\theta_s := \hat{u}_s^2 > 1/2$. Since $\hat{u}$ is a unit vector, it follows that if $\theta_s > 1/2$, then for any $q \neq s$, $\theta_q < 1/2$.

Consider the task of recognizing the distribution P_s. With the same assumptions on the environment (and the interaction between the environment and the learner), the role of the learner now is to recognize P_s, or equivalently to recognize the index s. From above, we obtain a simple reduction from this task to the task of subspace learning. Namely, given the output $\hat{\Pi} = \hat{u}\hat{u}^\top$ of $\mathcal{A}$, return the index s which maximizes the expression θ_s. We conclude from above that a successful subspace learner can be used to identify the distribution, i.e., to identify the index s.

B.2.2 $\Omega\left(\frac{d^2}{r^2\epsilon}\right)$ are needed to identify the index s

Assuming that the size of training sequence m is at most $o\left(\frac{d^2}{r^2\epsilon}\right)$, it is easily verified that with a non-negligible probability, the products observed by $\mathcal{A}$ algorithm are all zero. It follows that (with a non-negligible probability), $\mathcal{A}$ can not distinguish between the distributions P_s, and hence, $L(\hat{\Pi}) - L(\Pi^\star) > c\epsilon/2 > \epsilon$.

B.3 Proof of part 6 of Theorem 1

Proof 6 *Recall that $k = 1$, and $G = 1$. Recall the definition of the distributions $\{P_s : s \in [d]\}$. Let $\epsilon \in (0, 1/4)$. The zero vector is drawn with probability $1 - c\epsilon$ (for $c > 2$), and the vector e_s is drawn with probability $c\epsilon$. Note that $\mathbb{E}[xx^\top] = c\epsilon\, e_s e_s^\top$, and thus the optimal projection is simply $\Pi^\star = e_s e_s^\top = E_{s,s}$.*

Similarly to the previous subsections, we show that a successful subspace learner must identify the distribution. Let P_s be a concrete distribution. Denote by $\hat{\Pi} = \hat{u}\hat{u}^\top$ the output of the learner. Define $\theta_s = u_s^2$. As before, if follows that if $L(\hat{\Pi}) - L(\Pi^\star) < \epsilon < c\epsilon/2$, then $\theta_s > \theta_q$ for any $q \neq s$. That is, a successful subspace learner must identify the index s.

It is easily seen that if the size of the sample is at most $o\left(\frac{d}{r\epsilon}\right)$, then with a non-negligible probability, all the observations made by the learner are equal to zero. It follows that there exists a distribution P_j for which (with a non-negligible probability), $L(\hat{\Pi}) - L(\Pi^\star) > c\epsilon/2 > \epsilon$.

C Lower Bound on Subspace Learning in The Full Information Setting

As mentioned above, an upper bound of $O(G^2k/\epsilon^2)$ on the sample complexity of subspace learning has been proved by [4]. We now prove the optimality of this bound.

Theorem 4 *Assume that $k \leqslant d/2$. The sample complexity of subspace learning with full information over the domain $\mathcal{X} := \{x \in \mathbb{R}^d : \|x\|_2^2 \leqslant G, \|x\|_\infty^2 \leqslant 1\}$ is bounded by*

$$m(d, k, r = d, G, \epsilon) = \Omega(G^2k/\epsilon^2)$$

Hence,

$$m(d, k, r = d, G, \epsilon) \in \Theta(G^2k/\epsilon^2)\ .$$

We first sketch the idea behind the proof of Theorem 4. Let $\mathcal{U} = \{u_j\}_{j=1}^d$ be an orthogonal basis, such that $\|u_j\|^2 = G$ for each $j \in [d]$. A distribution over $\mathcal{U}$

is defined as follows. First, we draw a sequence $(b_1, \ldots, b_k) \in \{-1, 1\}^k$ uniformly at random. We associate the pair $\{u_i, u_{i+k}\}$ with a Bernoulli random variable B_i (coin) with parameter $p_i = \frac{1+b_i\alpha}{2}$, where α is proportional to ϵ/G. On each round, an instance x is drawn using the following procedure: First, an integer $i \in [k]$ is chosen uniformly at random. Second, the i-th coin is flipped (according to p_i), and x is set to be u_i or u_{i+k} accordingly. We show (Lemma 3) that a successful subspace learner "knows" the bias of most of the coins. Thus, we can reduce k independent tasks of coin identification to the task of subspace learning. A well-known result in statistics tells us that $\Omega(1/\alpha^2)$ samples are needed to identify a coin with bias α. The stated lower bound now follows.

Remark 2 *The concept of "free matrix lunch" (FML) is attributed to Manfred Warmuth ([18, 11]). A free matrix lunch is obtained whenever the "diagonal case", in which the distribution is supported on the standard basis is not "easier" (in terms of sample complexity) than the general case. In other words, a free matrix lunch holds when the learner does not earn anything (in the worst case) by knowing the eigensystem.*

Interestingly, the bound in Theorem 4 holds also when the eigensystem is restricted to be the standard eigensystem, i.e., when the orthogonal basis discussed above is the set $\{e_j : j \in [d]\}$. Hence, for the task of subspace learning in the full information setting, we obtain a free matrix lunch result.

C.1 Proof of Theorem 4

C.1.1 The adversarial distribution

We first construct a family of distributions which will pose a challenge to any subspace learner.

Example 1 *Assume that $k \leqslant d/2$. Let $\mathcal{U} = \{u_1, \ldots, u_{2k}\} \subseteq \mathbb{R}^d$ be $2k$ orthogonal vectors, such that $\|u_j\|^2 = G$ for each $j \in [d]$. Let $b_1, \ldots, b_k$ be independent Rademacher random variables (which are drawn in advance), and let $\alpha \in (0, 1)$ be a parameter. We define a distribution P over $\{u_1, \ldots, u_{2k}\}$ as follows:*

$$P(\{u_j\}) = \begin{cases} \frac{1+b_j\alpha}{2} & j \in [k] \\ \frac{1-b_j\alpha}{2} & j \in [2k]\backslash[k] \end{cases}$$

In words, the process of generating x can be described as follows:

- *The bias of k coins is randomly (and independently) determined to be $\frac{1\pm\alpha}{2}$.*
- *At each time, a coin $i \in [k]$ is chosen uniformly at random.*

- *The identity of x is chosen randomly to be u_i or u_{i+k}, according to the relevant coin.*

Example (1) defines a family $\mathcal{F}$ of distributions. The actual underlying distribution is chosen from $\mathcal{F}$ at random (by setting the bits $b_1, \ldots, b_k$). Note that proving a lower bound which holds in expectation over choosing $P \in \mathcal{F}$ at random implies that the lower bounds holds for some distribution in $\mathcal{F}$. This is a simple and very useful application of the *probabilistic method*.

C.1.2 A successful subspace learner is also a successful coins identifier

As we sketched before, a key ingredient of our analysis of the lower bounds is the relation between subspace learning and "coin identification". This relation is formalized next. We first need to introduce some additional notation. Consider the adversarial distributions described in Example (1). To specify the distribution P, fix some vector $(b_1, \ldots, b_k) \in \{-1, 1\}^k$. Let $\hat{\Pi} = \sum_{i=1}^k \hat{u}_i \hat{u}_i^\top \in \mathcal{P}_k^d$ (where $\hat{u}_1, \ldots, \hat{u}_k$ are orthonormal). For every $(i, j) \in [k] \times [2k]$, define $\theta_{i,j} = \frac{|\langle \hat{u}_i, u_j \rangle|}{\|\hat{u}_i\|_2 \|u_j\|_2} = \frac{|\langle \hat{u}_i, u_j \rangle|}{\sqrt{G}}$. Next, for each $j \in [2k]$, define $\theta_j = \sum_{i=1}^k \theta_{i,j}^2$. Also, define the set

$$J = \{j \in [k] : b_j = 1 \wedge \theta_j > \theta_{j+k}\} \cup \{j \in [k] : b_j = -1 \wedge \theta_j < \theta_{j+k}\} .$$

Finally, let $\beta = |J|/k$. For reasons that will become apparent shortly, we name J the set of identified coins. In this context, β is the fraction of identified coins.

Lemma 3 *Consider the adversarial distributions described in Example (1). Fix some vector $(b_1, \ldots, b_k) \in \{-1, 1\}^k$ to specify a distribution. Let $\hat{\Pi} = \sum_{i=1}^k \hat{u}_i \hat{u}_i^\top \in \mathcal{P}_k^d$ (where $\hat{u}_1, \ldots, \hat{u}_k$ are orthonormal). Assume that, for some $\epsilon \in (0, 1)$, $L(\hat{\Pi}) - \min_{\Pi \in \mathcal{P}_k^d} L(\Pi) < \epsilon$. Then, $\beta > 1 - \frac{2\epsilon}{\alpha G}$.*

Proof 7 *Assume w.l.o.g. that $b_1 = \ldots = b_k = 1$. Note that*

$$\mathbb{E}[xx^\top] = \frac{1}{k} \sum_{j=1}^k \left(\frac{1+\alpha}{2} u_j u_j^\top + \frac{1-\alpha}{2} u_{j+k} u_{j+k}^\top \right) .$$

The optimal projection is obtained by picking the largest eigenvectors of $\mathbb{E}[xx^\top]$. Precisely, the optimal projection matrix is $\Pi^\star = \sum_{i=1}^k (u_i/\|u_i\|)(u_i/\|u_i\|)^\top$. We

assume the loss as formulated in Equation (1). Let us calculate the loss of $\hat{\Pi}$*:*

$$\begin{aligned}
-L(\hat{\Pi}) &= \langle \sum_{i=1}^{k} \hat{u}_i \hat{u}_i^\top, \mathbb{E}[xx^\top] \rangle \\
&= \frac{1}{k} \sum_{i=1}^{k} \sum_{j=1}^{k} \left(\frac{1+\alpha}{2} \langle \hat{u}_i \hat{u}_i^\top, u_j u_j^\top \rangle + \frac{1-\alpha}{2} \langle \hat{u}_i \hat{u}_i^\top, u_{j+k} u_{j+k}^\top \rangle \right) \\
&= \frac{1}{k} \sum_{i=1}^{k} \sum_{j=1}^{k} \left(\frac{(1+\alpha) G \theta_{i,j}^2}{2} + \frac{(1-\alpha) G \theta_{i,j+k}^2}{2} \right) \\
&= \frac{1}{k} \sum_{j=1}^{k} \left(\frac{(1+\alpha) G \theta_j}{2} + \frac{(1-\alpha) G \theta_{j+k}}{2} \right) \\
&= \frac{G}{2k} \sum_{j=1}^{2k} \theta_j + \frac{\alpha G}{2k} \sum_{j=1}^{k} (\theta_j - \theta_{j+k}) \; . \qquad (8)
\end{aligned}$$

Note that the set $\left\{ \frac{u_j}{\|u_j\|} \right\}_{j=1}^{2k}$ *is orthonormal. Hence, for each* $i \in [k]$, $\sum_{j=1}^{2k} \theta_{i,j}^2 \leqslant 1$*. Hence,* $\sum_{j=1}^{2k} \theta_j \leqslant k$*, so the left-hand side of Equation (8) is at most* $\frac{G}{2}$*. This value is attained by* $\Pi^\star$*, and for simplicity, we will assume that is attained by* $\hat{\Pi}$ *as well. For the same reasons, for each* $i \in [k]$, $\sum_{j=1}^{k} (\theta_{i,j}^2 - \theta_{i,j+k}^2) \leqslant 1$*. Hence,* $\sum_{j=1}^{2k} (\theta_j - \theta_j) \leqslant k$*, so the right-hand size of Equation (8) is at most* $\alpha G/2$*. Once again, this value is attained by* $\Pi^\star$*. Our last observation is that if the* j*-th coin is not identified, i.e.,* $j \notin C$*, then* $\theta_j - \theta_{j+k} \leqslant 0$*. Combining these observations, we obtain that*

$$L(\hat{\Pi}) - L(\Pi^\star) \geqslant \alpha G (1-\beta)/2 \; .$$

The proof is completed by combining the assumption that $L(\hat{\Pi}) - L(\Pi^\star) < \epsilon$*.*

C.1.3 Lower bound on coin identification

We previously informally argued that $\Omega(1/\alpha^2)$ samples are needed to identify a coin with bias $\frac{1+\alpha}{2}$. If we have k such independent coins, then $\Omega(k/\alpha^2)$ are needed. Let us formalize this result.

A coin identification problem with parameter α is defined as follows. Consider a binary classification problem with a domain $[k]$ and label set $\{0,1\}$. The hypothesis class is the set $\mathcal{H} = \{0,1\}^{[k]}$. The underlying distribution over $[k] \times \{0,1\}$ is chosen at random using the following mechanism. The marginal distribution over $[k]$ is uniform, and the conditional probability over the label (coin) is determined

by $P(y=1|x=j)=\frac{1+b_j\alpha}{2}$, where each b_j is an independent Rademacher random variable (drawn in advance). We observe that this distribution is identical to the distribution defined (over a shattered set of size k) in [1, Theorem 5.2].

As usual, an algorithm for coin identification obtains an i.i.d. training sequence S according to P, and finally returns a hypothesis $h \in \mathcal{H}$. The generalization error of $h \in \mathcal{H}$, denoted $\mathrm{err}(h)$, is defined to be the probability that is misclassifies a new generated point (x,y), i.e., $\mathrm{err}(h) = P(h(x) \neq y)$. The next theorem follows from [1, Theorem 5.2].

Theorem 5 *Let $\mathcal{B}$ be an algorithm for coin identification, and let $\tilde{\epsilon} \in (0, 1/64)$. Consider a coin identification problem with $\alpha = 8\tilde{\epsilon}$. If $m < \frac{k}{8\tilde{\epsilon}^2}$, then there exists a distribution P for which*

$$\mathbb{E}_{S\sim P^m}\, \mathit{err}(\mathcal{B}(S)) - \min_{h\in\mathcal{H}} \mathit{err}(h) > \tilde{\epsilon}\ .$$

C.1.4 Concluding the theorem

We are in position to complete the reduction from the task of coin identification to the task of subspace learning, and consequently, conclude Theorem 4.

Let $G > 0$, $k \leqslant d/2$, and assume that $\epsilon \in (0, G/128)$. Let $\mathcal{A}$ be a subspace learner whose sample complexity is $m(d,k,r=d,G,\epsilon)$. We describe an algorithm $\mathcal{B}$ for coin identification (with parameter α) which uses $\mathcal{A}$ as a subroutine. We need to provide a map between the input of $\mathcal{B}$ to the input of $\mathcal{A}$. These inputs have the form of training sequences, which we will denote by S_{subspace} and S_{coins}, respectively. For each $j \in [k]$, the instance $(j,1)$ is associated with u_j, and the instance $(j,0)$ is associated with u_{j+k}. Clearly, the sequence provided to $\mathcal{A}$ is generated according to Example (1). To complete the reduction, we specify how the output of $\mathcal{B}$ is determined using the output of $\mathcal{A}$. This is straightforward. For each $j \in [k]$, the output hypothesis returned by $\mathcal{B}$ is defined by

$$h(j) = \begin{cases} 1 & \theta_j > \theta_{j+k} \\ 0 & \text{otherwise} \end{cases}\ .$$

Denote the fraction of identified coins by β (as in Lemma 3). Observe that each coin that is not identified, adds α to the relative error of $\mathcal{B}$. It follows from Lemma 3 that

$$\begin{aligned} \mathrm{err}(\mathcal{B}(S_{\mathrm{coins}})) - \min_{h\in\mathcal{H}} \mathrm{err}(h) &= \alpha(1-\beta) \\ &\leqslant 2\cdot \frac{L(\mathcal{A}(S_{\mathrm{subspace}})) - \min_{\Pi\in\mathcal{P}_k^d} L(\Pi)}{G} \end{aligned} \tag{9}$$

If $m \geqslant m(d, k, r = d, G, \epsilon)$, then the right-hand side of Equation (9) is at most $2\epsilon/G$. The proof is completed by applying Theorem 5 (with $\tilde{\epsilon} = 2\epsilon/G$).

D Decomposing elements in $\mathcal{C}_k^d$ into a convex combination of elements from $\mathcal{P}_k^d$

In this part we detail the decomposition procedure mentioned in Lemma 1. Recall that if $A \in \mathbb{S}^d$, then $\lambda(A)_1 \geqslant \ldots \geqslant \lambda(A)_d$ is the sequence of its eigenvalues. For convenience, we define two subroutines. Recall that $\mathrm{sort}(x, s)$ returns a set of s indices, corresponding to the s largest values of x. Second, we define a function $\mathrm{diag} : \mathbb{R}^d \to \mathbb{R}^{d \times d}$, which given a vector $x \in \mathbb{R}^d$, returns a diagonal matrix X with $X_{j,j} = x_j$.

Algorithm 4 Decomposition Procedure

Input: $W \in \mathcal{C}_k^d = \mathrm{conv}(\mathcal{P}_k^d)$
Output: $A \in \mathcal{P}_k^d$ s.t. $\mathbb{E}[A|W] = W$
perform eigendecomposition $W = U\mathrm{diag}(\lambda(A))U^\top$
$\lambda := k^{-1}(\lambda(A)_1, \ldots, \lambda(A)_d)$
$i := 1$
repeat
  $J := \mathrm{sort}(\lambda, k)$
  $c_i := k^{-1} \sum_{j \in J} e_j$
  $s := \min_{j \in J} \lambda_j$
  $l := \max_{j \in [d] \setminus J} \lambda_j$
  $\alpha_i := \min \left\{ sk, \sum_{j=1}^d \lambda_j - lk \right\}$
  $\lambda := \lambda - \alpha_i c_i$
  $\Pi_i := U\mathrm{diag}(kc_i)U^\top$
  $i := i + 1$
until $\lambda = (0, \ldots, 0)$
For each $j \in [i-1]$ choose $A = \Pi_j$ with probability α_j

It has been proved by [19] that the loop inside Algorithm 4 repeats at most d times, and decomposes W into a convex combination $\sum \alpha_i \Pi_i$ of elements from $\mathcal{P}_k^d$.